IA168 Algorithmic Game Theory

Tomáš Brázdil

Organization of This Course

Sources:

- Lectures (slides, notes)
  - based on several sources
  - slides are prepared for lectures, some stuff on greenboard ($\Rightarrow$ attend the lectures)
- Books:
  - Nisan/Roughgarden/Tardos/Vazirani, **Algorithmic Game Theory**, Cambridge University, 2007.
    Available online for free:
    http://www.cambridge.org/journals/nisan/downloads/Nisan_Non-printable.pdf
  - Tadelis, **Game Theory: An Introduction**, Princeton University Press, 2013

(I use various resources, so please, attend the lectures)

Evaluation

- **Oral exam**
- Homework

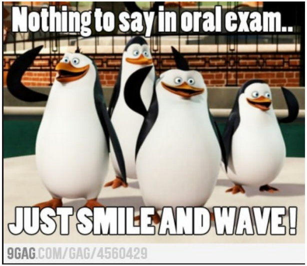

- 3 homework assignments

Notable features of the course

- No computer games course!
- **Very demanding!**
- Mathematical!

An unusual exam system!

You can repeat the oral exam as many times as needed (only the best grade goes into IS).

An example of an instruction email (from another course with the same system):

```
It is typically not sufficient to devote a single
afternoon to the preparation for the exam.
You have to know _everything_ (which means every
single thing) starting with the slide 42
and ending with the slide 245 with notable exceptions
of slides: 121 - 123, 137 - 140, 165, 167.
Proofs presented on the whiteboard are also mandatory.
```

Most importantly,

The previous slide is not a joke!

What is Algorithmic Game Theory?

First, what is the game theory?

According to the Oxford dictionary it is "the branch of mathematics concerned with the analysis of strategies for dealing with competitive situations where the outcome of a participant's choice of action depends critically on the actions of other participants"

According to Myerson it is "the study of mathematical models of conflict and cooperation between intelligent rational decision-makers"

What does the "algorithmic" mean?

- It means that we are "concerned with the computational questions that arise in game theory, and that enlighten game theory. In particular, questions about finding efficient algorithms to 'solve' games."

Let's have a look at some examples ....

Prisoner’s Dilemma

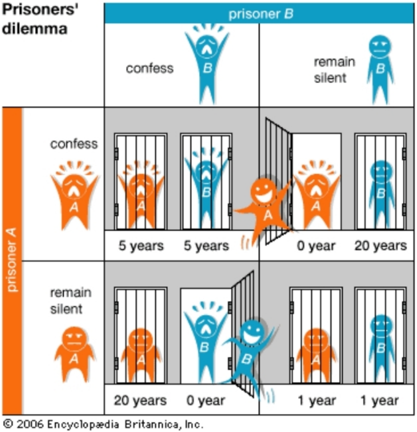

- Two suspects of a serious crime are arrested and imprisoned.
- Police has enough evidence of only petty theft, and to nail the suspects for the serious crime they need testimony from at least one of them.
- The suspects are interrogated separately without any possibility of communication.
- Each of the suspects is offered a deal: If he confesses (C) to the crime, he is free to go. The alternative is not to confess, that is remain silent (S).

Sentence depends on the behavior of both suspects.

The problem: What would the suspects do?

Prisoner's Dilemma – Solution(?)

	C	S
C	$-5, -5$	$0, -20$
S	$-20, 0$	$-1, -1$

Rational "row" suspect (or his adviser) may reason as follows:

- If my colleague chooses C, then playing C gives me -5 and playing S gives -20.
- If my colleague chooses S, then playing C gives me 0 and playing S gives -1.

In both cases C is clearly better (it *strictly dominates* the other strategy). If the other suspect's reasoning is the same, both choose C and get 5 years sentence.

Where is the dilemma? There is a solution (S, S) which is better for both players but needs some "central" authority to control the players.

Are there always "dominant" strategies?

Nash equilibria – Battle of Sexes

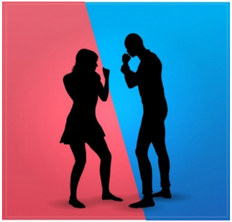

- A couple agreed to meet this evening, but cannot recall if they will be attending the opera or a football match.
- One of them wants to go to the football game. The other one to the opera. Both would prefer to go to the same place rather than different ones.

If they cannot communicate, where should they go?

Nash equilibria – Battle of Sexes

Battle of Sexes can be modeled as a game of two players (the couple) with the following payoffs:

	O	F
O	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

Apparently, no strategy of any player is dominant. A “solution”?

Note that whenever *both* players play O, then neither of them wants to *unilaterally* deviate from his strategy!

(O, O) is an example of a *Nash equilibrium* (as is (F, F))

Mixed Equilibria – Rock-Paper-Scissors

	R	P	S
R	$0, 0$	$-1, 1$	$1, -1$
P	$1, -1$	$0, 0$	$-1, 1$
S	$-1, 1$	$1, -1$	$0, 0$

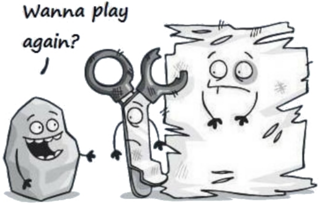

- This is an example of *zero-sum* games: whatever one of the players wins, the other one looses.
- What is an optimal behavior here? Is there a Nash equilibrium?

  Use *mixed strategies*: Each player plays each pure strategy with probability 1/3. The expected payoff of each player is 0 (even if one of the players changes his strategy, he still gets 0!).

Philosophical Issues in Games

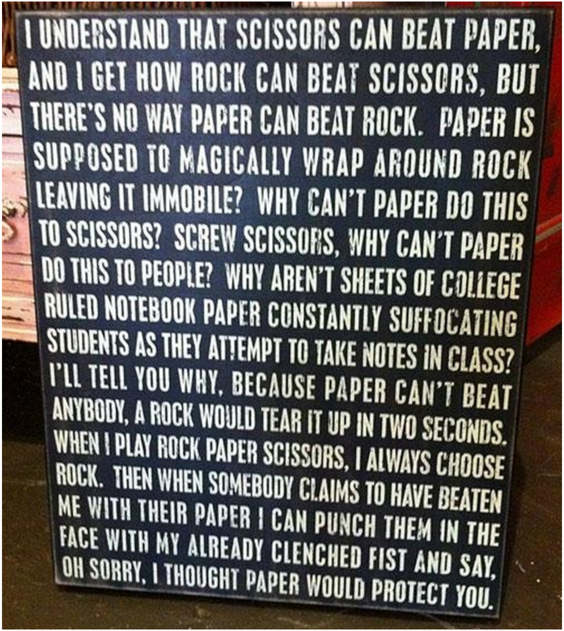

Dynamic Games

So far we have seen games in *strategic form* that are unable to capture games that unfold over time (such as chess).

For such purpose we need to use *extensive form* games:

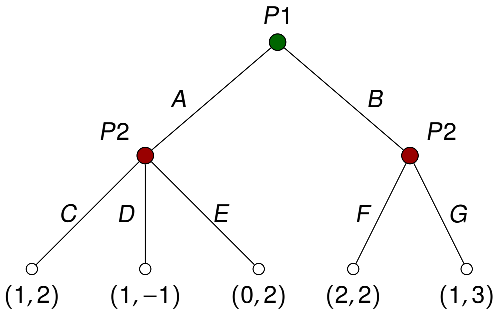

How to "solve" such games?

What is their relationship to the strategic form games?

Chance and Imperfect Information

Some decisions in the game tree may be by chance and controlled by neither player (e.g. Poker, Backgammon, etc.)

Sometimes a player may not be able to distinguish between several "positions" because he does not know all the information in them (Think a card game with opponent's cards hidden).

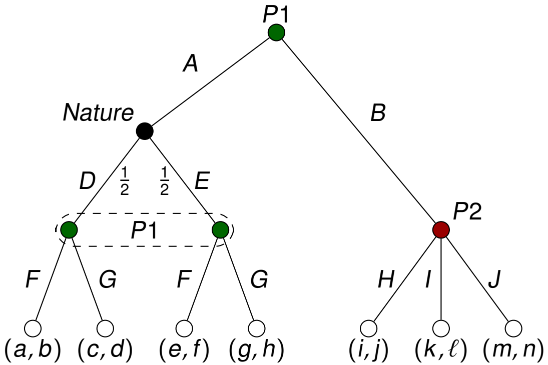

Again, how to solve such games?

Games of Incomplete Information

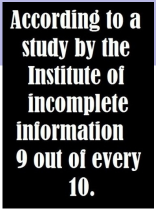

In all previous games the players knew all details of the game they played, and this fact was a “common knowledge”. This is not always the case.

Example: Sealed Bid Auction

- Two bidders are trying to purchase the same item.
- The bidders simultaneously submit bids b_1 and b_2 and the item is sold to the highest bidder at his bid price (first price auction)
- The payoff of the player 1 (and similarly for player 2) is calculated by

$$u_1(b_1, b_2) = \begin{cases} v_1 - b_1 & b_1 > b_2 \\ \frac{1}{2}(v_1 - b_1) & b_1 = b_2 \\ 0 & b_1 < b_2 \end{cases}$$

Here v_1 is the private value that player 1 assigns to the item and so the player 2 **does not know** u_1.

How to deal with such a game? Assume the “worst” private value? What if we have a partial knowledge about the private values?

Inefficiency of Equilibria

In Prisoner's Dilemma, the selfish behavior of suspects (the Nash equilibrium) results in somewhat worse than ideal situation.

	C	S
C	$-5, -5$	$0, -20$
S	$-20, 0$	$-1, -1$

Defining a *welfare function* W which to every pair of strategies assigns the sum of payoffs, we get $W(C, C) = -10$ but $W(S, S) = -2$.

The ratio $\frac{W(C,C)}{W(S,S)} = 5$ measures the inefficiency of "selfish-behavior" (C, C) w.r.t. the optimal “centralized” solution.

Price of Anarchy is the maximum ratio between values of equilibria and the value of an optimal solution.

Inefficiency of Equilibria – Selfish Routing

Consider a transportation system where many agents are trying to get from some initial location to a destination. Consider the welfare to be the average time for an agent to reach the destination. There are two versions:

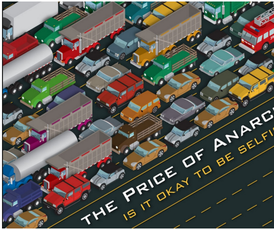

- "Centralized": A central authority tells each agent where to go.
- "Decentralized": Each agent selfishly minimizes his travel time.

Price of Anarchy measure the ratio between average travel time in these two cases.

Problem: Bound the price of anarchy over all routing games?

Games in Computer Science

Game theory is a core foundation of mathematical economics. But what does it have to do with CS?

- Games in AI: modeling of "rational" agents and their interactions.
- Games in machine learning: Generative adversarial networks, reinforcement learning
- Games in Algorithms: several game theoretic problems have a very interesting algorithmic status and are solved by interesting algorithms
- Games in modeling and analysis of reactive systems: program inputs viewed "adversarially", bisimulation games, etc.
- Games in computational complexity: Many complexity classes are definable in terms of games: PSPACE, polynomial hierarchy, etc.
- Games in Logic: modal and temporal logics, Ehrenfeucht-Fraisse games, etc.

Games in Computer Science

Games, the Internet and E-commerce: An extremely active research area at the intersection of CS and Economics

Basic idea: “The internet is a HUGE experiment in interaction between agents (both human and automated)”

How do we set up the rules of this game to harness “socially optimal” results?

Summary and Brief Overview

This is a *theoretical* course aimed at some fundamental results of game theory, often related to computer science

- We start with strategic form games (such as the Prisoner's dilemma), investigate several solution concepts (dominance, equilibria) and related algorithms.
- Then we consider repeated games which allow players to learn from history and/or to react to deviations of the other players.
- Subsequently, we move on to incomplete information games and auctions.
- Finally, we consider (in)efficiency of equilibria (such as the Price of Anarchy) and its properties on important classes of routing and network formation games.
- Remaining time will be devoted to selected topics from extensive form games, games on graphs etc.

Static Games of Complete Information

Strategic-Form Games

Solution concepts

Static Games of Complete Information – Intuition

Proceed in two steps:

1. Players *simultaneously and independently* choose their *strategies*. This means that players play without observing strategies chosen by other players.
2. Conditional on the players' strategies, *payoffs* are distributed to all players.

Complete information means that the following is *common knowledge* among players:

- all possible strategies of all players,
- what payoff is assigned to each combination of strategies.

Definition 1

A fact E is a *common knowledge* among players $\{1, \ldots, n\}$ if for every sequence $i_1, \ldots, i_k \in \{1, \ldots, n\}$ we have that i_1 knows that i_2 knows that ... i_{k-1} knows that i_k knows E.

The goal of each player is to maximize his payoff (and this fact is a common knowledge).

Strategic-Form Games

To formally represent static games of complete information we define *strategic-form games*.

Definition 2

A game in *strategic-form* (or normal-form) is an ordered triple $G = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$, in which:

- $N = \{1, 2, \ldots, n\}$ is a finite set of *players*.
- S_i is a set of *(pure) strategies* of player i, for every $i \in N$.

  A *strategy profile* is a vector of strategies of all players $(s_1, \ldots, s_n) \in S_1 \times \cdots \times S_n$.

  We denote the set of all strategy profiles by $S = S_1 \times \cdots \times S_n$.
- $u_i : S \to \mathbb{R}$ is a function associating each strategy profile $s = (s_1, \ldots, s_n) \in S$ with the *payoff* $u_i(s)$ to player i, for every player $i \in N$.

Definition 3

A *zero-sum* game G is one in which for all $s = (s_1, \ldots, s_n) \in S$ we have $u_1(s) + u_2(s) + \cdots + u_n(s) = 0$.

Example: Prisoner's Dilemma

- $N = \{1, 2\}$
- $S_1 = S_2 = \{S, C\}$
- u_1, u_2 are defined as follows:
  - $u_1(C, C) = -5$, $u_1(C, S) = 0$, $u_1(S, C) = -20$, $u_1(S, S) = -1$
  - $u_2(C, C) = -5$, $u_2(C, S) = -20$, $u_2(S, C) = 0$, $u_2(S, S) = -1$

(Is it zero sum?)

We usually write payoffs in the following form:

	C	S
C	$-5, -5$	$0, -20$
S	$-20, 0$	$-1, -1$

or as two matrices:

	C	S
C	-5	0
S	-20	-1

	C	S
C	-5	-20
S	0	-1

Example: Cournot Duopoly

- Two identical firms, players 1 and 2, produce some good. Denote by q_1 and q_2 quantities produced by firms 1 and 2, resp.
- The total quantity of products in the market is $q_1 + q_2$.
- The price of each item is $\kappa - q_1 - q_2$ (here κ is a positive constant)
- Firms 1 and 2 have per item production costs c_1 and c_2, resp.

Question: How these firms are going to behave?

We may model the situation using a strategic-form game.

Strategic-form game model $(N, (S_i)_{i \in N}, (u_i)_{i \in N})$

- $N = \{1, 2\}$
- $S_i = [0, \infty)$
- $u_1(q_1, q_2) = q_1(\kappa - q_1 - q_2) - q_1 c_1$
  $u_2(q_1, q_2) = q_2(\kappa - q_1 - q_2) - q_2 c_2$

Solution Concepts

A *solution concept* is a method of analyzing games with the objective of restricting the set of *all possible outcomes* to those that are *more reasonable than others.*

We will use term *equilibrium* for any one of the strategy profiles that emerges as one of the solution concepts' predictions.
(I follow the approach of Steven Tadelis here, it is not completely standard)

Example 4

Nash equilibrium is a solution concept. That is, we "solve" games by finding Nash equilibria and declare them to be reasonable outcomes.

Assumptions

Throughout the lecture we assume that:

1. Players are **rational**: a *rational* player is one who chooses his strategy to maximize his payoff.
2. Players are **intelligent**: An *intelligent* player knows everything about the game (actions and payoffs) and can make any inferences about the situation that we can make.
3. **Common knowledge**: The fact that players are rational and intelligent is a common knowledge among them.
4. **Self-enforcement**: Any prediction (or equilibrium) of a solution concept must be *self-enforcing*.

Here 4. implies non-cooperative game theory: Each player is in control of his actions, and he will stick to an action only if he finds it to be in his best interest.

Evaluating Solution Concepts

In order to evaluate our theory as a methodological tool we use the following criteria:

1. **Existence** (i.e., how often does it apply?): Solution concept should apply to a wide variety of games.

   E.g. We shall see that mixed Nash equilibria exist in all two player finite strategic-form games.

2. **Uniqueness** (How much does it restrict behavior?): We demand our solution concept to restrict the behavior as much as possible.

   E.g. So called strictly dominant strategy equilibria are always unique as opposed to Nash eq.

Solution Concepts – Pure Strategies

We will consider the following solution concepts:

- strict dominant strategy equilibrium
- iterated elimination of strictly dominated strategies (IESDS)
- rationalizability
- Nash equilibria

For now, let us concentrate on

pure strategies only!

I.e., no mixed strategies are allowed. We will generalize to mixed setting later.

Notation

- Let $N = \{1, \dots, n\}$ be a finite set and for each $i \in N$ let X_i be a set. Let $X := \prod_{i \in N} X_i = \{(x_1, \dots, x_n) \mid x_j \in X_j, j \in N\}$.
  - For $i \in N$ we define $X_{-i} := \prod_{j \neq i} X_j$, i.e.,

    $$X_{-i} = \{(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n) \mid x_j \in X_j, \forall j \neq i\}$$

  - An element of X_{-i} will be denoted by

    $$x_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$$

    We slightly abuse notation and write (x_i, x_{-i}) to denote $(x_1, \dots, x_i, \dots, x_n) \in X$.

Strict Dominance in Pure Strategies

Definition 5

Let $s_i, s'_i \in S_i$ be strategies of player i. Then s'_i is *strictly dominated* by s_i (write $s_i \succ s'_i$) if for any possible profile of the other players' strategies, $s_{-i} \in S_{-i}$, we have

$$u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i}) \quad \text{for all } s_{-i} \in S_{-i}$$

Is there a strictly dominated strategy in the Prisoner's dilemma?

	C	S
C	$-5, -5$	$0, -20$
S	$-20, 0$	$-1, -1$

Claim 1

An intelligent and rational player will never play a strictly dominated strategy.

Clearly, intelligence implies that the player should recognize dominated strategies, rationality implies that the player will avoid playing them.

Strictly Dominant Strategy Equilibrium in Pure Str.

Definition 6

$s_i \in S_i$ is *strictly dominant* if every other pure strategy of player i is strictly dominated by s_i.

Observe that every player has at most one strictly dominant strategy, and that strictly dominant strategies do not have to exist.

Claim 2

Any rational player will play the strictly dominant strategy (if it exists).

Definition 7

A strategy profile $s \in S$ is a *strictly dominant strategy equilibrium* if $s_i \in S_i$ is strictly dominant for all $i \in N$.

Corollary 8

If the strictly dominant strategy equilibrium exists, it is unique and rational players will play it.

Examples

In the Prisoner's dilemma:

	C	S
C	$-5,-5$	$0,-20$
S	$-20,0$	$-1,-1$

(C, C) is the strictly dominant strategy equilibrium.

In the Battle of Sexes:

	O	F
O	$2,1$	$0,0$
F	$0,0$	$1,2$

no strictly dominant strategies exist.

Indiana Jones and the Last Crusade

(Taken from Dixit & Nalebuff's "The Art of Strategy" and a lecture of Robert Marks)

Indiana Jones, his father, and the Nazis have all converged at the site of the Holy Grail. The two Joneses refuse to help the Nazis reach the last step. So the Nazis shoot Indiana's dad. Only the healing power of the Holy Grail can save the senior Dr. Jones from his mortal wound. Suitably motivated, Indiana leads the way to the Holy Grail. But there is one final challenge. He must choose between literally scores of chalices, only one of which is the cup of Christ. While the right cup brings eternal life, the wrong choice is fatal. The Nazi leader impatiently chooses a beautiful gold chalice, drinks the holy water, and dies from the sudden death that follows from the wrong choice. Indiana picks a wooden chalice, the cup of a carpenter. Exclaiming "There's only one way to find out" he dips the chalice into the font and drinks what he hopes is the cup of life. Upon discovering that he has chosen wisely, Indiana brings the cup to his father and the water heals the mortal wound.

Indiana Jones and the Last Crusade (cont.)

Indy Goofed

- Although this scene adds excitement, it is somewhat embarrassing that such a distinguished professor as Dr. Indiana Jones would overlook his dominant strategy.
- He should have given the water to his father without testing it first.
  - If Indiana has chosen the right cup, his father is still saved.
  - If Indiana has chosen the wrong cup, then his father dies but Indiana is spared.
- Testing the cup before giving it to his father doesn't help, since if Indiana has made the wrong choice, there is no second chance – Indiana dies from the water and his father dies from the wound.

Iterated Strict Dominance in Pure Strategies

We know that no rational player ever plays strictly dominated strategies.

As each player knows that each player is rational, each player knows that his opponents will not play strictly dominated strategies, and thus all opponents know that *effectively* they are facing a "smaller" game.

As rationality is common knowledge, everyone knows that everyone knows that the game is effectively smaller.

Thus, everyone knows that nobody will play strictly dominated strategies in the smaller game (and such strategies may indeed exist).

Because it is common knowledge that all players will perform this kind of reasoning again, the process can continue until no more strictly dominated strategies can be eliminated.

IESDS

The previous reasoning yields the **Iterated Elimination of Strictly Dominated Strategies (IESDS)**:

Define a sequence $D_i^0, D_i^1, D_i^2, \ldots$ of strategy sets of player i.
(Denote by G_{DS}^k the game obtained from G by restricting to D_i^k, $i \in N$.)

1. Initialize $k = 0$ and $D_i^0 = S_i$ for each $i \in N$.
2. For all players $i \in N$: Let D_i^{k+1} be the set of all pure strategies of D_i^k that are **not** strictly dominated in G_{DS}^k.
3. Let $k := k + 1$ and go to 2.

We say that $s_i \in S_i$ *survives IESDS* if $s_i \in D_i^k$ for all $k = 0, 1, 2, \ldots$

Definition 9

A strategy profile $s = (s_1, \ldots, s_n) \in S$ is an *IESDS equilibrium* if each s_i survives IESDS.

A game is *IESDS solvable* if it has a unique IESDS equilibrium.

Remark: If all S_i are *finite*, then in 2. we may remove only some of the strictly dominated strategies (not necessarily all). The result is *not* affected by the order of elimination since strictly dominated strategies remain strictly dominated even after removing some other strictly dominated strategies.

IESDS Examples

In the Prisoner's dilemma:

	C	S
C	$-5, -5$	$0, -20$
S	$-20, 0$	$-1, -1$

(C, C) is the only one surviving the first round of IESDS.

In the Battle of Sexes:

	O	F
O	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

all strategies survive all rounds (i.e. IESDS $\equiv$ anything may happen, sorry)

A Bit More Interesting Example

	L	C	R
L	$4, 3$	$5, 1$	$6, 2$
C	$2, 1$	$8, 4$	$3, 6$
R	$3, 0$	$9, 6$	$2, 8$

IESDS on greenboard!

Political Science Example

Hotelling (1929) and Downs (1957)

- $N = \{1, 2\}$
- $S_i = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$ (political and ideological spectrum)
- 10 voters belong to each position

  (Here 10 means ten percent in the real-world)
- Voters vote for the closest candidate. If there is a tie, then $\frac{1}{2}$ got to each candidate
- Payoff: The number of voters for the candidate; each candidate (selfishly) strives to maximize this number

Political Science Example

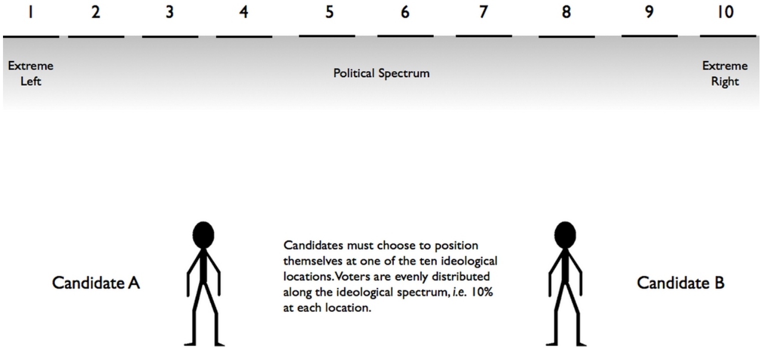

- 1 and 10 are the (only) strictly dominated strategies $\Rightarrow$ $D_1^1 = D_2^1 = \{2, \ldots, 9\}$
- in G_{DS}^1, 2 and 9 are the (only) strictly dominated strategies $\Rightarrow$ $D_1^2 = D_2^2 = \{3, \ldots, 8\}$
- $\ldots$
- only $5, 6$ survive IESDS

Belief & Best Response

IESDS eliminated apparently unreasonable behavior (leaving "reasonable" behavior implicitly untouched).

What if we rather want to actively preserve reasonable behavior? What is reasonable? .... what we believe is reasonable :-).

Intuition:

- Imagine that your colleague did something stupid
- What would you ask him? Usually, something like "What were you thinking?"
- The colleague may respond with a reasonable description of his *belief* in which his action was (one of) the best he could do

  (You may, of course, question the reasonableness of the belief)

Let us formalize this type of reasoning...

Belief & Best Response

Definition 10
A *belief* of player i is a pure strategy profile $s_{-i} \in S_{-i}$ of his opponents.

Definition 11
A strategy $s_i \in S_i$ of player i is a *best response* to a belief $s_{-i} \in S_{-i}$ if

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \quad \text{for all } s_i' \in S_i$$

Claim 3
A rational player who believes that his opponents will play $s_{-i} \in S_{-i}$ always chooses a best response to $s_{-i} \in S_{-i}$.

Definition 12
A strategy $s_i \in S_i$ is *never best response* if it is not a best response to any belief $s_{-i} \in S_{-i}$.

A rational player never plays any strategy that is never best response.

Best Response vs Strict Dominance

Proposition 1

If s_i is strictly dominated for player i, then it is never best response.

The opposite does not have to be true in pure strategies:

	X	Y
A	$1,1$	$1,1$
B	$2,1$	$0,1$
C	$0,1$	$2,1$

Here A is never best response but is strictly dominated neither by B, nor by C.

Elimination of Stupid Strategies = Rationalizability

Using similar iterated reasoning as for IESDS, strategies that are never best response can be iteratively eliminated.

Define a sequence $R_i^0, R_i^1, R_i^2, \ldots$ of strategy sets of player i.
(Denote by G_{Rat}^k the game obtained from G by restricting to R_i^k, $i \in N$.)

1. Initialize $k = 0$ and $R_i^0 = S_i$ for each $i \in N$.
2. For all players $i \in N$: Let R_i^{k+1} be the set of all strategies of R_i^k that are best responses to some beliefs in G_{Rat}^k.
3. Let $k := k + 1$ and go to 2.

We say that $s_i \in S_i$ is *rationalizable* if $s_i \in R_i^k$ for all $k = 0, 1, 2, \ldots$

Definition 13
A strategy profile $s = (s_1, \ldots, s_n) \in S$ is a *rationalizable equilibrium* if each s_i is rationalizable.

We say that a game is *solvable by rationalizability* if it has a unique rationalizable equilibrium.

(Warning: For some reasons, rationalizable strategies are almost always defined using mixed strategies!)

Rationalizability Examples

In the Prisoner's dilemma:

	C	S
C	$-5, -5$	$0, -20$
S	$-20, 0$	$-1, -1$

(C, C) is the only rationalizable equilibrium.

In the Battle of Sexes:

	O	F
O	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

all strategies are rationalizable.

Cournot Duopoly

$G = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$

- $N = \{1, 2\}$
- $S_i = [0, \infty)$
- $u_1(q_1, q_2) = q_1(\kappa - q_1 - q_2) - q_1 c_1 = (\kappa - c_1)q_1 - q_1^2 - q_1 q_2$
  $u_2(q_1, q_2) = q_2(\kappa - q_2 - q_1) - q_2 c_2 = (\kappa - c_2)q_2 - q_2^2 - q_2 q_1$

Assume for simplicity that $c_1 = c_2 = c$ and denote $\theta = \kappa - c$.

What is a best response of player 1 to a given q_2 ?

Solve $\frac{\delta u_1}{\delta q_1} = \theta - 2q_1 - q_2 = 0$, which gives that $q_1 = (\theta - q_2)/2$ is the only best response of player 1 to q_2.
Similarly, $q_2 = (\theta - q_1)/2$ is the only best response of player 2 to q_1.

Since $q_2 \geq 0$, we obtain that q_1 is never best response iff $q_1 > \theta/2$.
Similarly q_2 is never best response iff $q_2 > \theta/2$.

Thus $R_1^1 = R_2^1 = [0, \theta/2]$.

Cournot Duopoly

$G = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$

- $N = \{1, 2\}$
- $S_i = [0, \infty)$
- $u_1(q_1, q_2) = q_1(\kappa - q_1 - q_2) - q_1 c_1 = (\kappa - c_1)q_1 - q_1^2 - q_1 q_2$
  $u_2(q_1, q_2) = q_2(\kappa - q_2 - q_1) - q_2 c_2 = (\kappa - c_2)q_2 - q_2^2 - q_2 q_1$

Assume for simplicity that $c_1 = c_2 = c$ and denote $\theta = \kappa - c$.

Now, in G^1_{Rat}, we still have that $q_1 = (\theta - q_2)/2$ is the best response to q_2, and $q_2 = (\theta - q_1)/2$ the best resp. to q_1

Since $q_2 \in R_2^1 = [0, \theta/2]$, we obtain that q_1 is never best response iff $q_1 \in [0, \theta/4)$
Similarly q_2 is never best response iff $q_2 \in [0, \theta/4)$

Thus $R_1^2 = R_2^2 = [\theta/4, \theta/2]$.

....

Cournot Duopoly (cont.)

$G = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$

- $N = \{1, 2\}$
- $S_i = [0, \infty)$
- $u_1(q_1, q_2) = q_1(\kappa - q_1 - q_2) - q_1 c_1 = (\kappa - c_1)q_1 - q_1^2 - q_1 q_2$
  $u_2(q_1, q_2) = q_2(\kappa - q_2 - q_1) - q_2 c_2 = (\kappa - c_2)q_2 - q_2^2 - q_2 q_1$

Assume for simplicity that $c_1 = c_2 = c$ and denote $\theta = \kappa - c$.

In general, after $2k$ iterations we have $R_i^{2k} = R_i^{2k} = [\ell_k, r_k]$ where

- $r_k = (\theta - \ell_{k-1})/2$ for $k \geq 1$
- $\ell_k = (\theta - r_k)/2$ for $k \geq 1$ and $\ell_0 = 0$

Solving the recurrence we obtain

- $\ell_k = \theta/3 - \left(\frac{1}{4}\right)^k \theta/3$
- $r_k = \theta/3 + \left(\frac{1}{4}\right)^{k-1} \theta/6$

Hence, $\lim_{k\to\infty} \ell_k = \lim_{k\to\infty} r_k = \theta/3$ and thus $(\theta/3, \theta/3)$ is the only rationalizable equilibrium.

Cournot Duopoly (cont.)

$G = (N, (S_i)_{i\in N}, (u_i)_{i\in N})$

- $N = \{1, 2\}$
- $S_i = [0, \infty)$
- $u_1(q_1, q_2) = q_1(\kappa - q_1 - q_2) - q_1 c_1 = (\kappa - c_1)q_1 - q_1^2 - q_1 q_2$
  $u_2(q_1, q_2) = q_2(\kappa - q_2 - q_1) - q_2 c_2 = (\kappa - c_2)q_2 - q_2^2 - q_2 q_1$

Assume for simplicity that $c_1 = c_2 = c$ and denote $\theta = \kappa - c$.

Are $q_i = \theta/3$ the best outcomes possible? NO!

$$u_1(\theta/3, \theta/3) = u_2(\theta/3, \theta/3) = \theta^2/9$$

but

$$u_1(\theta/4, \theta/4) = u_2(\theta/4, \theta/4) = \theta^2/8$$

IESDS vs Rationalizability in Pure Strategies

Theorem 14

Assume that S is finite. Then for all k we have that $R_i^k \subseteq D_i^k$. That is, in particular, all rationalizable strategies survive IESDS.

The opposite inclusion does not have to be true in pure strategies:

	X	Y
A	$1,1$	$1,1$
B	$2,1$	$0,1$
C	$0,1$	$2,1$

Recall that A is never best response but is strictly dominated by neither B, nor C. That is, A survives IESDS but is not rationalizable.

Proof of Theorem 14

Claim

If s_i is a best response to s_{-i} in G^k_{Rat}, then s_i is a best response to s_{-i} in G.

Proof of the Claim. By induction on k. For $k = 0$ we have $G^k_{Rat} = G^0_{Rat} = G$ and the claim holds trivially.

Assume that the claim is true for some k and that s_i is a best response to s_{-i} in G^{k+1}_{Rat}. Let s'_i be a best response to s_{-i} in G^k_{Rat}.
Then $s'_i \in G^{k+1}_{Rat}$ since s'_i is *not* eliminated from G^k_{Rat}.
However, since s_i is a best response to s_{-i} in G^{k+1}_{Rat}, we get $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$.
Thus s_i is a best response to s_{-i} in G^k_{Rat}.

By induction hypothesis, s_i is a best response to s_{-i} in G and the claim has been proved.

Proof of Theorem 14

Keep in mind: If s_i is a best response to s_{-i} in G^k_{Rat}, then s_i is a best response to s_{-i} in G.

Now we prove $R^k_i \subseteq D^k_i$ for all players i by induction on k.
For $k = 0$ we have that $R^0_i = S_i = D^0_i$ by definition.
Assume that $R^k_i \subseteq D^k_i$ for some $k \geq 0$ and prove that $R^{k+1}_i \subseteq D^{k+1}_i$.
Let $s_i \in R^{k+1}_i$. Then there must be $s_{-i} \in R^k_{-i}$ such that

s_i is a best response to s_{-i} in G^k_{Rat}

(This follows from the fact that s_i has not been eliminated in G^k_{Rat}.)
By the claim, s_i is a best response to s_{-i} in G as well!
By induction hypothesis, $s_i \in R^{k+1}_i \subseteq R^k_i \subseteq D^k_i$ and $s_{-i} \in R^k_{-i} \subseteq D^k_{-i}$.
However, then s_i is a best response to s_{-i} in G^k_{DS}.
(This follows from the fact that the "best response" relationship of s_i and s_{-i} is preserved by removing arbitrarily many other strategies.)
Thus s_i is not strictly dominated in G^k_{DS} and $s_i \in D^{k+1}_i$. $\square$

Pinning Down Beliefs – Nash Equilibria

Criticism of previous approaches:

- Strictly dominant strategy equilibria often do not exist
- IESDS and rationalizability may not remove any strategies

Typical example is Battle of Sexes:

	O	F
O	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

Here all strategies are equally reasonable according to the above concepts.

But are all strategy profiles really equally reasonable?

Pinning Down Beliefs – Nash Equilibria

	O	F
O	$2,1$	$0,0$
F	$0,0$	$1,2$

Assume that each player has a belief about strategies of other players.

By Claim 3, each player plays a best response to his beliefs.

Is (O,F) as reasonable as (O,O) in this respect?

Note that if player 1 believes that player 2 plays O, then playing O is reasonable, and if player 2 believes that player 1 plays F, then playing F is reasonable. But such **beliefs cannot be correct together**!

(O,O) can be obtained as a profile where each player plays the best response to his belief and the **beliefs are correct**.

Nash Equilibrium

Nash equilibrium can be defined as a set of beliefs (one for each player) and a strategy profile in which every player plays a best response to his belief and each strategy of each player is consistent with beliefs of his opponents.

A usual definition is following:

Definition 15

A pure-strategy profile $s^* = (s_1^*, \ldots, s_n^*) \in S$ is a (pure) Nash equilibrium if s_i^* is a best response to s_{-i}^* for each $i \in N$, that is

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*) \quad \text{for all } s_i \in S_i \text{ and all } i \in N$$

Note that this definition is equivalent to the previous one in the sense that s_{-i}^* may be considered as the (consistent) belief of player i to which he plays a best response s_i^*

Nash Equilibria Examples

In the Prisoner's dilemma:

	C	S
C	$-5, -5$	$0, -20$
S	$-20, 0$	$-1, -1$

(C, C) is the only Nash equilibrium.

In the Battle of Sexes:

	O	F
O	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

only (O, O) and (F, F) are Nash equilibria.

In Cournot Duopoly, $(\theta/3, \theta/3)$ is the only Nash equilibrium.

(Best response relations: $q_1 = (\theta - q_2)/2$ and $q_2 = (\theta - q_1)/2$ are both satisfied only by $q_1 = q_2 = \theta/3$)

Example: Stag Hunt

Story:

- Two (in some versions more than two) hunters, players 1 and 2, can each choose to hunt

  - stag (S) = a large tasty meal
  - hare (H) = also tasty but small

- Hunting stag is much more demanding and forces of both players need to be joined (hare can be hunted individually)

Strategy-form game model: $N = \{1, 2\}$, $S_1 = S_2 = \{S, H\}$, the payoff:

	S	H
S	$5, 5$	$0, 3$
H	$3, 0$	$3, 3$

Two NE: (S, S), and (H, H), where the former is strictly better for each player than the latter! Which one is more reasonable?

Example: Stag Hunt

Strategy-form game model: $N = \{1, 2\}$, $S_1 = S_2 = \{S, H\}$, the payoff:

	S	H
S	$5, 5$	$0, 3$
H	$3, 0$	$3, 3$

Two NE: (S, S), and (H, H), where the former is strictly better for each player than the latter! Which one is more reasonable?

If each player believes that the other one will go for hare, then (H, H) is a reasonable outcome $\Rightarrow$ a society of individualists who do not cooperate at all.

If each player believes that the other will cooperate, then this anticipation is self-fulfilling and results in what can be called a cooperative society.

This is supposed to explain that in real world there are societies that have similar endowments, access to technology and physical environment but have very different achievements, all because of self-fulfilling beliefs (or *norms* of behavior).

Example: Stag Hunt

Strategy-form game model: $N = \{1, 2\}$, $S_1 = S_2 = \{S, H\}$, the payoff:

	S	H
S	$5, 5$	$0, 3$
H	$3, 0$	$3, 3$

Two NE: (S, S), and (H, H), where the former is strictly better for each player than the latter! Which one is more reasonable?

Another point of view: (H, H) is less risky

Minimum secured by playing S is 0 as opposed to 3 by playing H
(We will get to this *minimax* principle later)

So it seems to be rational to expect (H, H) (?)

Nash Equilibria vs Previous Concepts

Theorem 16

1. *If s^* is a strictly dominant strategy equilibrium, then it is the unique Nash equilibrium.*
2. *Each Nash equilibrium is rationalizable and survives IESDS.*
3. *If S is finite, neither rationalizability, nor IESDS creates new Nash equilibria.*

Proof: Homework!

Corollary 17

Assume that S is finite. If rationalizability or IESDS result in a unique strategy profile, then this profile is a Nash equilibrium.

Interpretations of Nash Equilibria

Except the two definitions, usual interpretations are following:

- When the goal is to give advice to all of the players in a game (i.e., to advise each player what strategy to choose), any advice that was not an equilibrium would have the unsettling property that there would always be some player for whom the advice was bad, in the sense that, if all other players followed the parts of the advice directed to them, it would be better for some player to do differently than he was advised. If the advice is an equilibrium, however, this will not be the case, because the advice to each player is the best response to the advice given to the other players.
- When the goal is prediction rather than prescription, a Nash equilibrium can also be interpreted as a potential stable point of a dynamic adjustment process in which individuals adjust their behavior to that of the other players in the game, searching for strategy choices that will give them better results.

Static Games of Complete Information

Mixed Strategies

Let's Mix It

As pointed out before, neither of the solution concepts has to exist in pure strategies

Example: Rock-Paper-sCissors

	R	P	C
R	$0, 0$	$-1, 1$	$1, -1$
P	$1, -1$	$0, 0$	$-1, 1$
C	$-1, 1$	$1, -1$	$0, 0$

There are no strictly dominant pure strategies

No strategy is strictly dominated (IESDS removes nothing)

Each strategy is a best response to some strategy of the opponent (rationalizability removes nothing)

No pure Nash equilibria: No *pure* strategy profile allows each player to play a best response to the strategy of the other player

How to solve this?

Let the players randomize their choice of pure strategies ....

Probability Distributions

Definition 18

Let A be a finite set. A *probability distribution over A* is a function $\sigma : A \rightarrow [0,1]$ such that $\sum_{a \in A} \sigma(a) = 1$.

We denote by $\Delta(A)$ the set of all probability distributions over A.

Example 19

Consider $A = \{a, b, c\}$ and a function $\sigma : A \rightarrow [0,1]$ such that $\sigma(a) = \frac{1}{4}$, $\sigma(b) = \frac{3}{4}$, and $\sigma(c) = 0$. Then $\sigma \in \Delta(A)$.

Mixed Strategies

Let us fix a strategic-form game $G = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$.

From now on, **assume two players and both S_i finite!**

$G = (\{1, 2\}, (S_1, S_2), (u_1, u_2))$

Definition 20

A *mixed strategy* of player i is a probability distribution $\sigma \in \Delta(S_i)$ over S_i. We denote by $\Sigma_i = \Delta(S_i)$ the set of all mixed strategies of player i.

We define $\Sigma := \Sigma_1 \times \Sigma_2$, the set of all *mixed strategy profiles*.

We identify each $s_i \in S_i$ with a mixed strategy σ that assigns probability one to s_i (and zero to other pure strategies).

For example, in rock-paper-scissors, the pure strategy R corresponds to σ_i which satisfies $\sigma_i(X) = \begin{cases} 1 & X = R \\ 0 & \text{otherwise} \end{cases}$

Mixed Strategy Profiles

Let $\sigma = (\sigma_1, \sigma_2)$ be a mixed strategy profile.

Intuitively, we assume that each player i *randomly* selects his pure strategy according to σ_i and *independently* of his opponents.

Thus for $s = (s_1, s_2) \in S = S_1 \times S_2$ we have that

$$\sigma(s) := \sigma_1(s_1) \cdot \sigma_2(s_2)$$

is the probability that the players randomly select the pure strategy profile s according to the mixed strategy profile σ.

(We abuse notation a bit here: σ denotes two things, a vector of mixed strategies as well as a probability distribution on S)

Mixed Strategies – Example

	R	P	C
R	$0,0$	$-1,1$	$1,-1$
P	$1,-1$	$0,0$	$-1,1$
C	$-1,1$	$1,-1$	$0,0$

An example of a mixed strategy σ_1: $\sigma_1(R) = \frac{1}{2}$, $\sigma_1(P) = \frac{1}{3}$, $\sigma_1(C) = \frac{1}{6}$.

Sometimes we write σ_1 as $(\frac{1}{2}(R), \frac{1}{3}(P), \frac{1}{6}(C))$, or only $(\frac{1}{2}, \frac{1}{3}, \frac{1}{6})$ if the order of pure strategies is fixed.

Consider a mixed strategy profile (σ_1, σ_2) where
$\sigma_1 = (\frac{1}{2}(R), \frac{1}{3}(P), \frac{1}{6}(C))$ and $\sigma_2 = (\frac{1}{3}(R), \frac{2}{3}(P), 0(C))$.
Then the probability $\sigma(R, P)$ that the pure strategy profile (R, P) will be played by players playing the mixed profile (σ_1, σ_2) is

$$\sigma_1(R) \cdot \sigma_2(P) = \frac{1}{2} \cdot \frac{2}{3} = \frac{1}{3}$$

Expected Payoff

... but now what is the suitable notion of payoff?

Definition 21

The *expected payoff* of player i under a mixed strategy profile $\sigma \in \Sigma$ is

$$u_i(\sigma) := \sum_{s \in S} \sigma(s) u_i(s) \qquad \left(= \sum_{s_1 \in S_1} \sum_{s_2 \in S_2} \sigma_1(s_1) \cdot \sigma_2(s_2) \cdot u_i(s_1, s_2) \right)$$

I.e., it is the "weighted average" of what player i wins under each pure strategy profile s, weighted by the probability of that profile.

Assumption: Every rational player strives to maximize his own expected payoff.

(This assumption is not always completely convincing ...)

Expected Payoff – Example

Matching Pennies:

	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

Each player secretly turns a penny to heads or tails, and then they reveal their choices simultaneously. If the pennies match, player 1 (row) wins, if they do not match, player 2 (column) wins.

Consider $\sigma_1 = (\frac{1}{3}(H), \frac{2}{3}(T))$ and $\sigma_2 = (\frac{1}{4}(H), \frac{3}{4}(T))$

$$\begin{aligned} u_1(\sigma_1, \sigma_2) &= \sum_{(X,Y)\in\{H,T\}^2} \sigma_1(X)\sigma_2(Y)u_1(X,Y) \\ &= \frac{1}{3}\frac{1}{4}1 + \frac{1}{3}\frac{3}{4}(-1) + \frac{2}{3}\frac{1}{4}(-1) + \frac{2}{3}\frac{3}{4}1 = \frac{1}{6} \end{aligned}$$

$$\begin{aligned} u_2(\sigma_1, \sigma_2) &= \sum_{(X,Y)\in\{H,T\}^2} \sigma_1(X)\sigma_2(Y)u_2(X,Y) \\ &= \frac{1}{3}\frac{1}{4}(-1) + \frac{1}{3}\frac{3}{4}1 + \frac{2}{3}\frac{1}{4}1 + \frac{2}{3}\frac{3}{4}(-1) = -\frac{1}{6} \end{aligned}$$

Solution Concepts

We revisit the following solution concepts in mixed strategies:

- strict dominant strategy equilibrium
- IESDS equilibrium
- rationalizable equilibria
- Nash equilibria

From now on, when I say a *strategy* I implicitly mean a

mixed strategy.

In order to deal with efficiency issues we assume that the size of the game G is defined by $|G| := |N| + \sum_{i \in N} |S_i| + \sum_{i \in N} |u_i|$ where $|u_i| = \sum_{s \in S} |u_i(s)|$ and $|u_i(s)|$ is the length of a binary encoding of $u_i(s)$ (we assume that rational numbers are encoded as quotients of two binary integers)

Note that, in particular, $|G| > |S|$.

Strict Dominance in Mixed Strategies

Definition 22
Let $\sigma_1, \sigma_1' \in \Sigma_1$ be (mixed) strategies of player 1. Then σ_1' is *strictly dominated* by σ_1 (write $\sigma_1' \prec \sigma_1$) if

$$u_1(\sigma_1, s_2) > u_1(\sigma_1', s_2) \quad \text{for all } s_2 \in S_2$$

(Symmetrically for player 2.)

Comment: The above condition is equivalent to

$$u_1(\sigma_1, \sigma_2) > u_1(\sigma_1', \sigma_2) \quad \text{for all strategies } \sigma_2 \in \Sigma_2$$

Strict Dominance in Mixed Strategies

Example 23

	X	Y
A	3	0
B	0	3
C	1	1

Is there a strictly dominated strategy?

Question: Is there a game with at least one strictly dominated strategy but without strictly dominated *pure* strategies?

Strictly Dominant Strategy Equilibrium

Definition 24
$\sigma_i \in \Sigma_i$ is *strictly dominant* if every other mixed strategy of player i is strictly dominated by σ_i.

Definition 25
A strategy profile $\sigma \in \Sigma$ is a *strictly dominant strategy equilibrium* if $\sigma_i \in \Sigma_i$ is strictly dominant for each $i \in N$.

Proposition 2
If the strictly dominant strategy equilibrium exists, it is unique; all its strategies are pure, and rational players will play it.

Proof.
Homework. □

To compute the strictly dominant strategy equilibrium, it is sufficient to consider only pure strategies.

IESDS in Mixed Strategies

Define a sequence $D_i^0, D_i^1, D_i^2, \ldots$ of strategy sets of player i.
(Denote by G_{DS}^k the game obtained from G by restricting the pure strategy sets to D_i^k, $i \in N$.)

1. Initialize $k = 0$ and $D_i^0 = S_i$ for each $i \in N$.
2. For all players $i \in N$: Let D_i^{k+1} be the set of all pure strategies of D_i^k that are *not* strictly dominated in G_{DS}^k by *mixed strategies*.
3. Let $k := k + 1$ and go to 2.

We say that $s_i \in S_i$ *survives IESDS* if $s_i \in D_i^k$ for all $k = 0, 1, 2, \ldots$

Definition 26
A strategy profile $s = (s_1, s_2) \in S$ is an *IESDS equilibrium* if both s_1 and s_2 survive IESDS.

Each D_i^{k+1} can be computed in polynomial time using *linear programming*.

IESDS in Mixed Strategie – Example

	X	Y
A	3	0
B	0	3
C	1	1

Let us have a look at the first iteration of IESDS.

Observe that A, B are not strictly dominated by any mixed strategy.

Let us construct a set of constraints on mixed strategies (possibly) strictly dominating C:

$$\begin{aligned} 3x_A + 0x_B + x_C &> 1 && \text{Row's payoff against } X \\ 0x_A + 3x_B + x_C &> 1 && \text{Row's payoff against } Y \\ x_A, x_B, x_C &\geq 0 \\ x_A + x_B + x_C &= 1 && \text{x's must make a distribution} \end{aligned}$$

How to solve this?

Intermezzo: Linear Programming

Linear programming is a technique for optimization of a linear objective function, subject to linear (non-strict) inequality constraints.

Formally, a linear program in so called *canonical form* looks like this:

$$\text{maximize} \sum_{j=1}^{m} c_j x_j \qquad \textit{(objective function)}$$

$$\begin{aligned} \text{subject to} \sum_{j=1}^{m} a_{ij} x_j &\le b_i \qquad 1 \le i \le n \\ x_j &\ge 0 \qquad 1 \le j \le m \end{aligned} \qquad \textit{(constraints)}$$

Here a_{ij}, b_k and c_j are real numbers and x_j's are real variables.

A *feasible solution* is an assignment of real numbers to the variables x_j, $1 \le j \le m$, so that the *constraints* are satisfied.

An *optimal solution* is a feasible solution which maximizes the *objective function* $\sum_{j=1}^{m} c_j x_j$.

Intermezzo: Complexity of Linear Programming

We assume that coefficients a_{ij}, b_k and c_j are encoded in binary (more precisely, as fractions of two integers encoded in binary).

Theorem 27 (Khachiyan, Doklady Akademii Nauk SSSR, 1979)

There is an algorithm which for any linear program computes an optimal solution in polynomial time.

The algorithm uses so called ellipsoid method.

In practice, the Khachiyan's is not used. Usually **simplex algorithm** is used even though its theoretical complexity is exponential.

There is also a polynomial time algorithm (by Karmarkar) which has better complexity upper bounds than the Khachiyan's and sometimes works even better than the simplex.

There exist several advanced linear programming solvers (usually parts of larger optimization packages) implementing various heuristics for solving large scale problems, sensitivity analysis, etc.

For more info see

http://en.wikipedia.org/wiki/Linear_programming#Solvers_and_scripting_.28programming.29_languages

IESDS in Mixed Strategie – Example

	X	Y
A	3	0
B	0	3
C	1	1

The linear program for deciding whether C is strictly dominated: The program maximizes y under the following constraints:

$$
\begin{aligned}
3x_A + 0x_B + x_C &\geq 1 + y && \text{Row's payoff against } X \\
0x_A + 3x_B + x_C &\geq 1 + y && \text{Row's payoff against } Y \\
x_A, x_B, x_C &\geq 0 && \\
x_A + x_B + x_C &= 1 && \text{x's must make a distribution} \\
y &\geq 0 &&
\end{aligned}
$$

Here y just implements the strict inequality using $\geq$, we look for a solution with $y > 0$.

The maximum $y = \frac{1}{2}$ is attained at $x_A = \frac{1}{2}$ and $x_B = \frac{1}{2}$.

IESDS – Algorithm

Note that in step 2 it is not sufficient to consider pure strategies. Consider the following zero sum game:

	X	Y
A	3	0
B	0	3
C	1	1

C is strictly dominated by $(\sigma_1(A), \sigma_1(B), \sigma_1(C)) = (\frac{1}{2}, \frac{1}{2}, 0)$ but no strategy is strictly dominated in pure strategies.

Best Response in Mixed Strategies

Definition 28
A *(mixed) belief* of player 1 is a mixed strategy σ_2 of player 2 (and vice versa).

Definition 29
$\sigma_1 \in \Sigma_1$ is a *best response* to a belief $\sigma_2 \in \Sigma_2$ if

$$u_1(\sigma_1, \sigma_2) \geq u_1(s_1, \sigma_2) \quad \text{for all } s_1 \in S_1$$

Denote by $BR_1(\sigma_2)$ the set of all best responses of player 1. (Symmetrically for player 2.)

Comment: The above condition is equivalent to

$$u_1(\sigma_1, \sigma_2) \geq u_1(\sigma_1', \sigma_2) \quad \text{for all } \sigma_1' \in \Sigma_1$$

Best Response – Example

Consider a game with the following payoffs of player 1:

	X	Y
A	2	0
B	0	2
C	1	1

- Player 1 (row) plays $\sigma_1 = (a(A), b(B), c(C))$.
- Player 2 (column) plays $(q(X), (1-q)(Y))$ (we write just q).

Compute $BR_1(q)$.

Rationalizability in Mixed Strategies (Two Players)

Assumption: *A rational player* 1 *with a belief* σ_2 *always plays a best response to* σ_2 *(the same for player 2).*

Definition 30

A pure strategy $s_1 \in S_1$ of player 1 is *never best response* if it is not a best response to any belief σ_2 (similarly for player 2).

No rational player plays a strategy that is never best response.

Rationalizability in Mixed Strategies (Two Players)

Define a sequence $R_i^0, R_i^1, R_i^2, \ldots$ of strategy sets of player i.
(Denote by G_{Rat}^k the game obtained from G by restricting the pure strategy sets to R_i^k, $i \in N$.)

1. Initialize $k = 0$ and $R_i^0 = S_i$ for each $i \in N$.
2. For all players $i \in N$: Let R_i^{k+1} be the set of all strategies of R_i^k that are *best responses to some (mixed) beliefs in* G_{Rat}^k.
3. Let $k := k + 1$ and go to 2.

We say that $s_i \in S_i$ is *rationalizable* if $s_i \in R_i^k$ for all $k = 0, 1, 2, \ldots$

Definition 31
A strategy profile $s = (s_1, s_2) \in S$ is a *rationalizable equilibrium* if both s_1 and s_2 are rationalizable.

Rationalizability vs IESDS (Two Players)

	X	Y
A	3	0
B	0	3
C	1	1

What pure strategies of player 1 are strictly dominated?

What pure strategies of player 1 are never best responses?

Observation: The set of strictly dominated pure strategies coincides with the set of pure never best responses!

... and this holds in general for two player games:

Theorem 32

A pure strategy s_1 of player 1 is never best response to any belief σ_2 **iff** *s_1 is strictly dominated by a strategy $\sigma_1 \in \Sigma_1$ (similarly for player 2).*

It follows that a strategy of S_i survives IESDS **iff** it is rationalizable.

Mixed Nash Equilibrium

Definition 33

A mixed-strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*) \in \Sigma$ is a (mixed) Nash equilibrium if σ_1^* is a best response to σ_2^* and σ_2^* is a best response to σ_1^*. That is

$$u_1(\sigma_1^*, \sigma_2^*) \geq u_1(s_1, \sigma_2^*) \quad \text{for all } s_1 \in S_1$$

$$u_2(\sigma_1^*, \sigma_2^*) \geq u_2(\sigma_1^*, s_2) \quad \text{for all } s_2 \in S_2$$

The above condition is equivalent to

$$u_1(\sigma_1^*, \sigma_2^*) \geq u_1(\sigma_1, \sigma_2^*) \quad \text{for all } \sigma_1 \in \Sigma_1$$
$$u_2(\sigma_1^*, \sigma_2^*) \geq u_2(\sigma_1^*, \sigma_2) \quad \text{for all } \sigma_2 \in \Sigma_2$$

Theorem 34 (Nash 1950)

Every finite game in strategic form has a Nash equilibrium.

This is THE fundamental theorem of game theory.

Example: Matching Pennies

	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

Player 1 (row) plays $(p(H), (1-p)(T))$ (we write just p) and player 2 (column) plays $(q(H), (1-q)(T))$ (we write q).

Compute all Nash equilibria.

What are the expected payoffs of playing pure strategies for player 1?

$$u_1(H, q) = 2q - 1 \text{ and } u_1(T, q) = 1 - 2q$$

Then
$u_1(p, q) = pu_1(H, q) + (1-p)u_1(T, q) = p(2q-1) + (1-p)(1-2q)$.

We obtain the best response correspondence BR_1:

$$BR_1(q) = \begin{cases} T & \text{if } q < \frac{1}{2} \\ p \in [0, 1] & \text{if } q = \frac{1}{2} \\ H & \text{if } q > \frac{1}{2} \end{cases}$$

Example: Matching Pennies

	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

Player 1 (row) plays $(p(H), (1-p)(T))$ (we write just p) and player 2 (column) plays $(q(H), (1-q)(T))$ (we write q).

Compute all Nash equilibria.

Similarly for player 2 :

$$u_2(p, H) = 1 - 2p \text{ and } u_2(p, T) = 2p - 1$$

$u_2(p, q) = qu_2(p, H) + (1-q)u_2(p, T) = q(1-2p) + (1-q)(2p-1)$
We obtain best-response relation BR_2:

$$BR_2(p) = \begin{cases} H & \text{if } p < \frac{1}{2} \\ q \in [0, 1] & \text{if } p = \frac{1}{2} \\ T & \text{if } p > \frac{1}{2} \end{cases}$$

The only "intersection" of BR_1 and BR_2 is the only Nash equilibrium $\sigma_1 = \sigma_2 = (\frac{1}{2}, \frac{1}{2})$.

Support Enumeration

Computing Mixed Nash Equilibria

Lemma 35
Every Nash equilibrium $\sigma^* = (\sigma_1^*, \sigma_2^*) \in \Sigma$ *satisfies*

- $u_1(s_1, \sigma_2^*) = u_1(\sigma^*)$ *for* $s_1 \in supp(\sigma_1^*)$
- $u_2(\sigma_1^*, s_2) = u_2(\sigma^*)$ *for* $s_2 \in supp(\sigma_2^*)$

Proof. W.l.o.g. consider only the player 1 and assume that σ^* is a Nash equilibrium.

The latter assumption implies $u_1(s_1, \sigma_2^*) \leq u_1(\sigma^*)$ for all $s_1 \in S_1$.

Now, if there exists $s_1' \in supp(\sigma_1^*) \subseteq S_1$ satisfying $u_1(s_1', \sigma_2^*) < u_1(\sigma^*)$, then because $\sigma_1^*(s_1') > 0$ we have

$$u_1(\sigma^*) = \sum_{s_1 \in S_1} \sigma_1^*(s_1) u_1(s_1, \sigma_2^*) < \sum_{s_1 \in S_1} \sigma_1^*(s_1) u_1(\sigma^*) = u_1(\sigma^*)$$

A contradiction.

Thus $u_1(s_1, \sigma_2^*) = u_1(\sigma^*)$ for all $s_1 \in supp(\sigma_1^*)$.

Example: Matching Pennies

	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

Player 1 (row) plays $(p(H), (1-p)(T))$ (we write just p) and player 2 (column) plays $(q(H), (1-q)(T))$ (we write q).

Compute all Nash equilibria.

There are no pure strategy equilibria.

There are no equilibria where only player 1 randomizes:

Indeed, assume that (p, H) is such an equilibrium. Then by Lemma 35,

$$1 = u_1(H, H) = u_1(T, H) = -1$$

a contradiction. Also, (p, T) cannot be an equilibrium.

Similarly, there is no NE where only player 2 randomizes.

Example: Matching Pennies

	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

Player 1 (row) plays $(p(H), (1-p)(T))$ (we write just p) and player 2 (column) plays $(q(H), (1-q)(T))$ (we write q).

Compute all Nash equilibria.

Assume that both players randomize, i.e., $p, q \in (0, 1)$.

The expected payoffs of playing pure strategies for player 1:

$$u_1(H, q) = 2q - 1 \text{ and } u_1(T, q) = 1 - 2q$$

Similarly for player 2 :

$$u_2(p, H) = 1 - 2p \text{ and } u_1(p, T) = 2p - 1$$

By Lemma 35, such Nash equilibria must satisfy:

$$2q - 1 = 1 - 2q \qquad \text{and} \qquad 1 - 2p = 2p - 1$$

That is $p = q = \frac{1}{2}$ is the only Nash equilibrium.

Example: Battle of Sexes

	O	F
O	$2,1$	$0,0$
F	$0,0$	$1,2$

Player 1 (row) plays $(p(O),(1-p)(F))$ (we write just p) and player 2 (column) plays $(q(O),(1-q)(F))$ (we write q).

Compute all Nash equilibria.

There are two pure strategy equilibria (O,O) and (F,F), no Nash equilibrium where only one player randomizes.

Now assume that

- player 1 (row) plays $(p(O),(1-p)(F))$ (we write just p) and
- player 2 (column) plays $(q(O),(1-q)(F))$ (we write q)

where $p,q \in (0,1)$.

By Lemma 35, such Nash equilibria must satisfy:

$$2q = 1-q \qquad \text{and} \qquad p = 2(1-p)$$

This holds only for $q = \frac{1}{3}$ and $p = \frac{2}{3}$.

An Algorithm?

What did we do in the previous examples?

We went through all support combinations for both players.
(pure, one player mixing, both mixing)

For each pair of supports we tried to find equilibria in strategies with these supports.
(in Battle of Sexes: two pure, no equilibrium with just one player mixing, one equilibrium when both mixing)

Whenever one of the *supports* was non-singleton, we reduced computation of Nash equilibria to *linear equations*.

Computing Mixed Nash Equilibria

Lemma 36

Let $\sigma^* = (\sigma_1^*, \sigma_2^*) \in \Sigma$ be a mixed profile. Assume that there exist $w_1, w_2 \in \mathbb{R}$ such that

- $u_1(s_1, \sigma_2^*) = w_1$ *for* $s_1 \in supp(\sigma_1^*)$
- $u_1(s_1, \sigma_2^*) \leq w_1$ *for* $s_1 \notin supp(\sigma_1^*)$
- $u_2(\sigma_1^*, s_2) = w_2$ *for* $s_2 \in supp(\sigma_2^*)$
- $u_2(\sigma_1^*, s_2) \leq w_2$ *for* $s_2 \notin supp(\sigma_2^*)$

Then $u_1(\sigma^*) = w_1$ and $u_2(\sigma^*) = w_2$, and σ^* is a Nash equilibrium.

Proof. Consider just the player 1 (for pl. 2 similarly):

$$u_1(\sigma^*) = \sum_{s_1 \in S_1} \sigma^*(s_1) u_1(s_1, \sigma_2^*) = \sum_{s_1 \in supp(\sigma_1^*)} \sigma^*(s_1) u_1(s_1, \sigma_2^*)$$

$$= \sum_{s_1 \in supp(\sigma_1^*)} \sigma^*(s_1) w_1 = w_1 \sum_{s_1 \in supp(\sigma_1^*)} \sigma^*(s_1) = w_1$$

Now the fact that σ^* is a Nash equilibrium follows from the definition.

How to Compute Mixed Nash Equilibria?

Every Nash equilibrium $\sigma^* = (\sigma_1^*, \sigma_2^*)$ can be computed by finding appropriate w_1, w_2 so that

- $u_1(s_1, \sigma_2^*) = w_1$ for $s_1 \in supp(\sigma_1^*)$
- $u_1(s_1, \sigma_2^*) \leq w_1$ for $s_1 \notin supp(\sigma_1^*)$
- $u_2(\sigma_1^*, s_2) = w_2$ for $s_2 \in supp(\sigma_2^*)$
- $u_2(\sigma_1^*, s_2) \leq w_2$ for $s_2 \notin supp(\sigma_2^*)$

Indeed,

- by Lemma 36, all σ^* and w_1, w_2 satisfying the above inequalities give a Nash equilibrium σ^* with $u_1(\sigma^*) = w_1$ and $u_2(\sigma^*) = w_2$,
- by Lemma 35, for every Nash equilibrium σ^* choosing $w_1 = u_1(\sigma^*)$ and $w_2 = u_2(\sigma^*)$ satisfies the above inequalities.

Suppose that we somehow know the supports $supp(\sigma_1^*), supp(\sigma_2^*)$ for some Nash equilibrium $\sigma^* = (\sigma_1^*, \sigma_2^*)$ (which itself is unknown to us).

We may consider all $\sigma_i^*(s_i)$'s and both w_1, w_2's as variables and use the above conditions to design a system of inequalities capturing Nash equilibria with the given support sets $supp(\sigma_1^*), supp(\sigma_2^*)$.

Support Enumeration

To simplify notation, assume that for every i we have $S_i = \{1, \dots, m_i\}$.
Then $\sigma_i(j)$ is the probability of the pure strategy j in the mixed strategy σ_i.

Fix supports $supp_i \subseteq S_i$ for every $i \in \{1, 2\}$ and consider the following system of constraints with variables $\sigma_1(1), \dots, \sigma_1(m_1), \sigma_2(1), \dots, \sigma_2(m_2), w_1, w_2$:

1. For all $k \in supp_1$ and all $\ell \in supp_2$:

$$\sum_{\ell' \in S_2} \sigma_2(\ell') u_1(k, \ell') = w_1 \qquad \sum_{k' \in S_1} \sigma_1(k') u_2(k', \ell) = w_2$$

2. For all $k \notin supp_1$ and all $\ell \notin supp_2$:

$$\sum_{\ell' \in S_2} \sigma_2(\ell') u_1(k, \ell') \leq w_1 \qquad \sum_{k' \in S_1} \sigma_1(k') u_2(k', \ell) \leq w_2$$

3. For all $i \in \{1, 2\}$: $\sigma_i(1) + \dots + \sigma_i(m_i) = 1$.
4. For all $i \in \{1, 2\}$ and all $k \in supp_i$: $\sigma_i(k) \geq 0$.
5. For all $i \in \{1, 2\}$ and all $k \notin supp_i$: $\sigma_i(k) = 0$.

Support Enumeration

The constraints are *linear* for two player games!

How to find $supp_1$ and $supp_2$? ... Just guess!

Input: A two-player strategic-form game G with strategy sets $S_1 = \{1, \ldots, m_1\}$ and $S_2 = \{1, \ldots, m_2\}$ and rational payoffs u_1, u_2.

Output: A Nash equilibrium σ^*.

Algorithm: For all possible $supp_1 \subseteq S_1$ and $supp_2 \subseteq S_2$:

- Check if the corresponding system of linear constraints (from the previous slide) has a feasible solution σ^*, w_1^*, w_2^*.
- If so, STOP: the feasible solution σ^* is a Nash equilibrium satisfying $u_i(\sigma^*) = w_i^*$.

Question: How many possible subsets $supp_1, supp_2$ are there to try?
Answer: $2^{(m_1+m_2)}$

So, unfortunately, the algorithm requires worst-case exponential time.

Remarks on Support Enumeration

- The algorithm combined with Theorem 34 and properties of linear programming imply that every finite two-player game has a rational Nash equilibrium (furthermore, the rational numbers have polynomial representation in binary).
- The algorithm can be used to compute *all* Nash equilibria.

  (There are algorithms for computing (a finite representation of) a set of all feasible solutions of a given linear constraint system.)
- The algorithm can be used to compute "good" equilibria.

  For example, to find a Nash equilibrium maximizing the sum of all expected payoffs (the "social welfare") it suffices to solve the system of constraints while maximizing $w_1 + w_2$. More precisely, the algorithm can be modified as follows:
  - Initialize $W := -\infty$ (W stores the current maximum welfare)
  - For all possible $supp_1 \subseteq S_1$ and $supp_2 \subseteq S_2$:
    - Find the maximum value $\max(w_1 + w_2)$ of $w_1 + w_2$ so that the constraints are satisfiable (using linear programming).
    - Put $W := \max\{W, \max(w_1 + w_2)\}$.
  - Return W.

Remarks on Support Enumeration (Cont.)

Similar trick works for any notion of "good" NE that can be expressed using a linear objective function and (additional) linear constraints in variables $\sigma_i(j)$ and w_i.

(e.g., maximize payoff of player 1, minimize payoff of player 2 and keep probability of playing the strategy 1 below $1/2$, etc.)

Complexity Results – (Two Players)

Theorem 37
Given a two-player game in strategic form, a mixed Nash equilibrium can be computed in exponential time.

Theorem 38
All the following problems are NP-complete: Given a two-player game in strategic form, does it have

1. *a NE in which player 1 has utility at least a given amount v ?*
2. *a NE in which the sum of expected payoffs of the two players is at least a given amount v ?*
3. *a NE with a support of size greater than a given number?*
4. *a NE whose support contains a given strategy s ?*
5. *a NE whose support does not contain a given strategy s ?*
6. *....*

NP-hardness can be proved using reduction from SAT.

The Reduction (It's Short and Sweet)

Definition 4 *Let ϕ be a Boolean formula in conjunctive normal form (representing a SAT instance). Let V be its set of variables (with $|V| = n$), L the set of corresponding literals (a positive and a negative one for each variable[6]), and C its set of clauses. The function $v : L \to V$ gives the variable corresponding to a literal,* e.g., *$v(x_1) = v(-x_1) = x_1$. We define $G_\epsilon(\phi)$ to be the following finite symmetric 2-player game in normal form. Let $\Sigma = \Sigma_1 = \Sigma_2 = L \cup V \cup C \cup \{f\}$. Let the utility functions be*

- $u_1(l^1, l^2) = u_2(l^2, l^1) = n - 1$ *for all* $l^1, l^2 \in L$ *with* $l^1 \neq -l^2$;
- $u_1(l, -l) = u_2(-l, l) = n - 4$ *for all* $l \in L$;
- $u_1(l, x) = u_2(x, l) = n - 4$ *for all* $l \in L$, $x \in \Sigma - L - \{f\}$;
- $u_1(v, l) = u_2(l, v) = n$ *for all* $v \in V$, $l \in L$ *with* $v(l) \neq v$;
- $u_1(v, l) = u_2(l, v) = 0$ *for all* $v \in V$, $l \in L$ *with* $v(l) = v$;
- $u_1(v, x) = u_2(x, v) = n - 4$ *for all* $v \in V$, $x \in \Sigma - L - \{f\}$;
- $u_1(c, l) = u_2(l, c) = n$ *for all* $c \in C$, $l \in L$ *with* $l \notin c$;
- $u_1(c, l) = u_2(l, c) = 0$ *for all* $c \in C$, $l \in L$ *with* $l \in c$;
- $u_1(c, x) = u_2(x, c) = n - 4$ *for all* $c \in C$, $x \in \Sigma - L - \{f\}$;
- $u_1(x, f) = u_2(f, x) = 0$ *for all* $x \in \Sigma - \{f\}$;
- $u_1(f, f) = u_2(f, f) = \epsilon$;
- $u_1(f, x) = u_2(x, f) = n - 1$ *for all* $x \in \Sigma - \{f\}$.

Theorem 1 *If $(l_1, l_2, \ldots, l_n)$ (where $v(l_i) = x_i$) satisfies ϕ, then there is a Nash equilibrium of $G_\epsilon(\phi)$ where both players play l_i with probability $\frac{1}{n}$, with expected utility $n - 1$ for each player. The only other Nash equilibrium is the one where both players play f, and receive expected utility ϵ each.*

... But What is The Exact Complexity of *Computing* Nash Equilibria in Two Player Games?

Let us concentrate on the problem of computing one Nash equilibrium (sometimes called the *sample equilibrium problem*).

As the class NP consists of decision problems, it cannot be directly used to characterize complexity of the sample equilibrium problem.

We use complexity classes of *function problems* such as FP, FNP, etc. The sample equilibrium problem belongs to the complexity class PPAD (which is a subclass of TFNP) for two-player games.

A binary relation P(x,y) is in TFNP if and only if there is a deterministic polynomial time algorithm that can determine whether P(x,y) holds given both x and y, and for every x, there exists a y which is at most polynomially longer than x such that P(x,y) holds.

Can we do better than FNP (i.e. exponential time)?

In what follows we show that the sample equilibrium problem can be solved in polynomial time for zero-sum two-player games.

(Using a beautiful characterization of all Nash equilibria)

MaxMin

Definition 39
$\sigma_1^* \in \Sigma_1$ is a *maxmin* strategy of player 1 if

$$\sigma_1^* \in \underset{\sigma_1 \in \Sigma_1}{\operatorname{argmax}} \min_{s_2 \in S_2} u_1(\sigma_1, s_2) \quad \left(= \underset{\sigma_1 \in \Sigma_1}{\operatorname{argmax}} \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2)\right)$$

(Intuitively, a *maxmin* strategy σ_1^* maximizes player 1's worst-case payoff in the situation where player 2 strives to cause the greatest harm to player 1.)

Similarly, $\sigma_2^* \in \Sigma_2$ is a *maxmin* strategy of player 2 if

$$\sigma_2^* \in \underset{\sigma_2 \in \Sigma_2}{\operatorname{argmax}} \min_{s_1 \in S_1} u_2(s_1, \sigma_2)$$

Which assuming zero-sum games, i.e. $u_1 = -u_2$, becomes

$$\sigma_2^* \in \underset{\sigma_2 \in \Sigma_2}{\operatorname{argmin}} \max_{s_1 \in S_1} u_1(s_1, \sigma_2) \quad \left(= \underset{\sigma_2 \in \Sigma_2}{\operatorname{argmin}} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2)\right)$$

Note the same payoff function for both players!!

Zero-Sum Games: von Neumann's Theorem

Theorem 40 (von Neumann)
Assume a two-player **zero-sum** *game. Then*

$$\max_{\sigma_1 \in \Sigma_1} \min_{s_2 \in S_2} u_1(\sigma_1, s_2) = \min_{\sigma_2 \in \Sigma_2} \max_{s \in S_1} u_1(s_1, \sigma_2)$$

Morever, $\sigma^* = (\sigma_1^*, \sigma_2^*) \in \Sigma$ *is a Nash equilibrium* **iff** *both* σ_1^* *and* σ_2^* *are maxmin.*

So to compute a Nash equilibrium it suffices to compute (arbitrary) maxmin strategies for both players.

Zero-Sum Two-Player Games – Computing NE

Assume $S_1 = \{1, \ldots, m_1\}$ and $S_2 = \{1, \ldots, m_2\}$.

We want to compute

$$\sigma_1^* \in \operatorname*{argmax}_{\sigma_1 \in \Sigma_1} \min_{\ell \in S_2} u_1(\sigma_1, \ell)$$

Consider a linear program with variables $\sigma_1(1), \ldots, \sigma_1(m_1), v$:

$$\begin{array}{rll}
\textbf{maximize:} & v & \\
\textbf{subject to:} & \displaystyle\sum_{k=1}^{m_1} \sigma_1(k) \cdot u_1(k, \ell) \geq v & \ell = 1, \ldots, m_2 \\
& \displaystyle\sum_{k=1}^{m_1} \sigma_1(k) = 1 & \\
& \sigma_1(k) \geq 0 & k = 1, \ldots, m_1
\end{array}$$

Lemma 41

$\sigma_1^* \in \operatorname{argmax}_{\sigma_1 \in \Sigma_1} \min_{\ell \in S_2} u_1(\sigma_1, \ell)$ **iff** *assigning* $\sigma_1(k) := \sigma_1^*(k)$ *and* $v := \min_{\ell \in S_2} u_1(\sigma_1^*, \ell)$ *gives an optimal solution.*

Summary:

- We have reduced computation of NE to computation of maxmin strategies for both players.
- Maxmin strategies can be computed using linear programming in polynomial time.
- That is, Nash equilibria in zero-sum two-player games can be computed in polynomial time.

Strategic-Form Games – Conclusion

We have considered *static games of complete information*, i.e., "one-shot" games where the players know exactly what game they are playing.
We modeled such games using *strategic-form games*.

We have considered both pure strategy setting and mixed strategy setting.

In both cases, we considered four solution concepts:

- Strictly dominant strategies
- Iterative elimination of strictly dominated strategies
- Rationalizability (i.e., iterative elimination of strategies that are never best responses)
- Nash equilibria

Strategic-Form Games – Conclusion

In pure strategy setting:

1. Strictly dominant strategy equilibrium survives IESDS, rationalizability and is the unique Nash equilibrium (if it exists)
2. In finite games, rationalizable equilibria survive IESDS, IESDS preserves the set of Nash equilibria
3. In finite games, rationalizability preserves Nash equilibria

In mixed setting:

1. In finite two player games, IESDS and rationalizability coincide.
2. Strictly dominant strategy equilibrium survives IESDS (rationalizability) and is the unique Nash equilibrium (if it exists)
3. In finite games, IESDS (rationalizability) preserves Nash equilibria

The proofs for 2. and 3. in the mixed setting are similar to corresponding proofs in the pure setting.

Algorithms

- Strictly dominant strategy equilibria coincide in pure and mixed settings, and can be computed in polynomial time.
- IESDS and rationalizability can be implemented in polynomial time in the pure setting as well as in the mixed setting

  In the mixed setting, linear programming is needed to implement one step of IESDS (rationalizability).
- Nash equilibria can be computed for two-player games
  - in polynomial time for zero-sum games
    (using von Neumann's theorem and linear programming)
  - in exponential time using support enumeration
  - in PPAD using Lemke-Howson (omitted)

Loose Ends – Modes of Dominance

To simplify, let us consider only **pure strategies**.

Let $s_i, s_i' \in S_i$. Then s_i' is *strictly dominated* by s_i if
$u_i(s_i, s_{-i}) > u_i(s_i', s_{-i}) \quad$ for all $s_{-i} \in S_{-i}$.

Let $s_i, s_i' \in S_i$. Then s_i' is *weakly dominated* by s_i if
$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$ for all $s_{-i} \in S_{-i}$ and there is $s'_{-i} \in S_{-i}$ such that $u_i(s_i, s'_{-i}) > u_i(s_i', s'_{-i})$.

Let $s_i, s_i' \in S_i$. Then s_i' is *very weakly dominated* by s_i if
$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$ for all $s_{-i} \in S_{-i}$.

A strategy is (strictly, weakly, very weakly) dominant if it (strictly, weakly, very weakly) dominates any other strategy.

Claim 4

Any pure strategy profile $s \in S$ such that each s_i is very weakly dominant is a Nash equilibrium.

The same claim can be proved in the mixed strategy setting.